Heidi Seibold, Achim Zeileis and Torsten Hothorn*

Model-Based Recursive Partitioning for Subgroup Analyses

DOI 10.1515/ijb-2015-0032

Abstract: The identification of patient subgroups with differential treatment effects is the first step towards individualised treatments. A current draft guideline by the EMA discusses potentials and problems in subgroup analyses and formulated challenges to the development of appropriate statistical procedures for the data-driven identification of patient subgroups. We introduce model-based recursive partitioning as a procedure for the automated detection of patient subgroups that are identifiable by predictive factors. The method starts with a model for the overall treatment effect as defined for the primary analysis in the study protocol and uses measures for detecting parameter instabilities in this treatment effect. The procedure produces a segmented model with differential treatment parameters corresponding to each patient subgroup. The subgroups are linked to predictive factors by means of a decision tree. The method is applied to the search for subgroups of patients suffering from amyotrophic lateral sclerosis that differ with respect to their Riluzole treatment effect, the only currently approved drug for this disease.

Keywords: subgroup analysis, personalized medicine, treatment efficacy, permutation test, amyotrophic lateral sclerosis

1 Introduction

With the rise of personalised medicine, the search for individual treatments poses challenges to the development of appropriate statistical methods. Subgroup analyses following a traditional statistical assessment of an overall treatment effect of a new therapy aim at identifying three groups of patients: (1) those who benefit from the new therapy, (2) those who do not benefit, and (3) those whose clinical outcome under the new therapy is worse than under alternative therapies. Such post-hoc subgroup analyses potentially lead to better benefit-risk decisions and treatment recommendations but are subject to all kind of biases and can hardly be performed under full statistical error control. Therefore, the European Medicines Agency (EMA) recently published a draft of a guideline for the investigation of subgroups in confirmatory clinical trials [1] that discusses potential areas of application, necessity, pitfalls, and good practice in subgroup analyses. In the guideline draft, three scenarios in which exploratory investigation of subgroups is of special interest were identified:

Scenario 1: “The clinical data presented are overall statistically persuasive with therapeutic efficacy demonstrated globally. It is of interest to verify that the conclusions of therapeutic efficacy (and safety) apply consistently across subgroups of the clinical trial population.”

Scenario 2: “The clinical data presented are overall statistically persuasive but with therapeutic efficacy or benefit/risk which is borderline or unconvincing and it is of interest to identify post-hoc a subgroup, where efficacy and risk-benefit is convincing.”

Scenario 3: “The clinical data presented fail to establish statistically persuasive evidence but there is interest in identifying a subgroup, where a relevant treatment effect and compelling evidence of a favourable risk-benefit profile can be assessed.”

*Corresponding author: **Torsten Hothorn**, Department of Biostatistics, Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Zurich, Switzerland

Heidi Seibold, Department of Biostatistics, Epidemiology, Biostatistics and Prevention Institute, University of Zurich, Zurich, Switzerland

Achim Zeileis, Department of Statistics, Faculty of Economics and Statistics University of Innsbruck, Innsbruck, Austria

Especially in trials with highly heterogeneous study populations, subgroup analyses can help to reduce the variability of the estimated overall treatment effect by splitting the study population into more homogeneous subgroups.

Information about the individual treatment effect might be available from cross-over trials or from counterfactual analyses of parallel-group designs [2, 3]. These individual effects can then be linked to potentially predictive variables. In the absence of such information, most importantly in the case of parallel-group designs studied here, subgroup analyses can be seen as the search for or specification of treatment $\times$ covariate interactions and we proceed along this path. A covariate measures a patient characteristic that potentially explains the patient's individual treatment effect. In the commonly applied models with linear predictors, such as the linear, generalised linear or linear transformation models, the specification of higher-order interaction terms and especially the subsequent inference are known to be burdensome. For non-categorical covariates, it is a priori unclear how one can derive a subgroup from a significant treatment $\times$ covariate interaction.

Automated interaction detection [4], today known as recursive partitioning methods or simply “trees”, was suggested as an interaction search procedure more than 50 years ago, and has had a very active development community ever since. Although the application of trees for subgroup identification seems to be straightforward, no generally applicable method is available [5]. The main technical problem is that classical trees were developed for identifying higher-order covariate interactions but additional work is required to restrict interactions to treatment $\times$ covariate interactions. Due to the non-parametric nature of most tree models, blending trees with the linear models typically used to describe the treatment effect is challenging.

While setting up such automated procedures for subgroup identification, one has to bear in mind that the impact of a covariate on the endpoint can be prognostic, predictive, or both. Prognostic factors have a direct impact on the endpoint, independent of the treatment applied. This corresponds to a main effect. A predictive factor explains a differential treatment effect, i.e. a treatment $\times$ covariate interaction term. Both the main and the treatment interaction terms are important for factors that are prognostic and predictive at the same time [6].

In our analysis, we aimed at detecting subgroups of patients suffering from amyotrophic lateral sclerosis (ALS) in which the subgroups differ in the effect of treatment with Riluzole, the only approved drug for ALS treatment today. The two endpoints of interest are a functional endpoint assessing the patient's ability to handle daily life and the overall survival time. We estimated the overall treatment effect of Riluzole using four different base models; the choice of the model depended on the measurement scale of the endpoint. A normal generalised linear model (GLM) with log-link was used for the sum-score of the functional endpoint, and item-specific proportional odds models were used for the decomposed score. For the right-censored survival times, we used a parametric Weibull model and a semiparametric Cox model. Our aim was to partition these linear models with respect to the treatment effect parameter and to develop a segmented model that includes treatment $\times$ covariate interactions that describe the relevant subgroups.

We applied model-based recursive partitioning [7] to the functional and survival models describing the effect of Riluzole on ALS patients in order to obtain subgroups with a differential treatment effect. The main advantage of embedding our subgroup analysis into this general framework of model partitioning is that one can partition the base model used for analysing the overall treatment effect, regardless of the measurement scale of the endpoint. The method allows us to focus attention on predictive factors, while other terms, such as the effects of strata or nuisance parameters, can be held fixed.

Section 2 introduces the general framework for subgroup identification and compares the new procedure to methods published previously in the light of this general theoretical framework. In Section 3, we present results of our subgroup analysis of Riluzole treatment of ALS patients and discuss the patient subgroups and corresponding differential treatment effects found.

2 Model-based recursive partitioning for subgroup identification

Subgroup analyses require the definition of a parameter describing the treatment effect. In clinical trials, this parameter is typically already contained in the model that was defined in the study protocol for the

analysis of the primary endpoint. The treatment effect was estimated in the primary analysis under the assumption that the corresponding parameter is universally applicable to all patients. In the presence of subgroups, this assumption does not hold and these patient subgroups differ in their treatment effect. If we assume that the different treatment effects can be understood as a function of patient characteristics, the patient subgroups can be identified by estimating this treatment effect function. Model-based recursive partitioning can be employed as a procedure for the estimation of such a treatment effect function and the identification of the corresponding patient subgroups. The name of the procedure comes from the nature of the algorithm that recursively partitions the initial model used for the analysis of the primary endpoint.

2.1 Model and algorithm

We started with a model $\mathcal{M}((Y, \mathbf{X}), \vartheta)$ that describes the conditional distribution of the primary endpoint Y (or certain characteristics of this distribution) as a function of the treatment arm and potentially further covariates (both contained in $\mathbf{X}$) through parameters ϑ as defined in the study protocol. The parameter vector $\vartheta = (\alpha, \beta, \gamma, \sigma)^\top$ typically contains one or more intercept parameters α , one or more treatment effect parameters β , other model parameters γ , e.g. effects of covariates, and potential nuisance parameters σ , e.g. the error variance in a linear model. The estimator is defined as the minimizer of an objective function Ψ , which usually is the negative log-likelihood:

$$\hat{\vartheta} = \arg \min_{\vartheta} \sum_{i=1}^N \Psi((y, \mathbf{x})_i, \vartheta). \quad (1)$$

Estimating ϑ is equivalent to solving the score equation

$$\sum_{i=1}^N \frac{\partial \Psi((y, \mathbf{x})_i, \vartheta)}{\partial \vartheta} = \sum_{i=1}^N \psi((y, \mathbf{x})_i, \vartheta) = 0, \quad (2)$$

where ψ is the score function, i.e. the gradient of the objective function Ψ with respect to ϑ . The model framework is more general than the log-likelihood framework because Ψ is not necessarily a negative log-likelihood function.

In the presence of patient subgroups that differ in their treatment effect β , an estimate $\hat{\beta}$ obtained for all patients $i=1, \dots, N$ in the study only reflects the mean treatment effect but ignores that the success or failure of a specific treatment might depend on additional characteristics of each individual patient. We describe patient subgroups as a partition $\{\mathcal{B}_b\}$ ($b=1, \dots, B$) of all patients $i=1, \dots, N$. The subgroup-specific model parameters are then $\vartheta(b)$. These parameters can in general be seen as varying coefficients [8], however they may depend on several patient characteristics and are always step functions with a different level for each subgroup and not only a smoothly varying coefficient for one single predictive variable.

Since we are searching for predictive and prognostic factors, we are only interested in subgroups that differ in the intercept or the treatment effect or both as explained in Section 2.2. With $\vartheta(b) = (\alpha(b), \beta(b), \gamma, \sigma)^\top$ we assume that the effects of covariates and nuisance parameters are constant for all patients. The partition $\{\mathcal{B}_b\}$ is defined by J partitioning variables $\mathbf{Z} = (Z_1, \dots, Z_J) \in \mathcal{Z}$; in other words, $\{\mathcal{B}_b\}$ is a hypercube in the J -dimensional sample space $\mathcal{Z}$. These partitioning variables $\mathbf{Z}$ are the additional patient characteristics that potentially influence $\alpha(b)$ and $\beta(b)$. If for example gender were a predictive factor in a given treatment-endpoint relationship, it would be a patient characteristic that is involved in forming the partitions. If the partition $\{\mathcal{B}_b\}$ is known, the partitioned model parameters $\vartheta(b)$ could be estimated by minimising the segmented objective function:

$$(\hat{\vartheta}(b))_{b=1, \dots, B} = \arg \min_{\vartheta(b)} \sum_{i=1}^N \sum_{b=1}^B \mathbb{1}(\mathbf{z}_i \in \mathcal{B}_b) \Psi((y, \mathbf{x})_i, \vartheta(b)), \quad (3)$$

where $\mathbb{1}$ denotes the indicator function and $(y, \mathbf{x})_i, \mathbf{z}_i$ are the realisations of $(Y, \mathbf{X})$ and $\mathbf{Z}$ for the i -th patient. This allows us to write the subgroup-specific intercept and treatment parameters as a function of the partitioning variables

$$\alpha(\mathbf{z}) = \sum_{b=1}^B \mathbb{1}(\mathbf{z} \in \mathcal{B}_b) \cdot \alpha(b) \quad \text{and} \quad \beta(\mathbf{z}) = \sum_{b=1}^B \mathbb{1}(\mathbf{z} \in \mathcal{B}_b) \cdot \beta(b).$$

Without any a priori knowledge about the partition $\{\mathcal{B}_b\}$, we want to estimate the functions $\alpha(\mathbf{z})$ and $\beta(\mathbf{z})$ by means of model-based recursive partitioning. The main idea underlying this method is the ability to detect parameter instabilities, i.e. non-constant parameters in a parametric or semiparametric model, by looking at the score function. Because we are only interested in detecting non-constant intercepts $\alpha(\mathbf{z})$ and treatment effects $\beta(\mathbf{z})$, we focus on the partial score functions $\psi_\alpha((Y, \mathbf{X}), \vartheta) = \partial\Psi((Y, \mathbf{X}), \vartheta)/\partial\alpha$ and $\psi_\beta((Y, \mathbf{X}), \vartheta) = \partial\Psi((Y, \mathbf{X}), \vartheta)/\partial\beta$. If the model parameters are in fact constant and do not depend on any of the partitioning variables $\mathbf{Z}$, the partial score functions $\psi_\alpha((Y, \mathbf{X}), \vartheta)$ and $\psi_\beta((Y, \mathbf{X}), \vartheta)$ are independent of $\mathbf{Z}$. Consequently, parameter instability corresponds to a correlation between either of the partial score functions and at least one of the partitioning variables $Z_1, \dots, Z_J$. In order to formally detect deviations from independence between the partial score functions and the partitioning variables, model-based recursive partitioning utilises independence tests. The null hypotheses

$$\begin{aligned} H_0^{\alpha,j}: \quad & \psi_\alpha((Y, \mathbf{X}), \hat{\vartheta}) \perp Z_j, j=1, \dots, J \\ & \text{and} \\ H_0^{\beta,j}: \quad & \psi_\beta((Y, \mathbf{X}), \hat{\vartheta}) \perp Z_j, j=1, \dots, J \end{aligned}$$

for a given model $\mathcal{M}((Y, \mathbf{X}), \hat{\vartheta})$ state that the partial score functions with respect to α and β , respectively, are independent of the partitioning variable Z_j ($j=1, \dots, J$). Hence, these null hypotheses correspond to an appropriate model fit regarding the intercept and treatment parameter. Because the partial score functions under the null hypotheses are at least asymptotically normal in many model families, asymptotic M-fluctuation tests with appropriate correction for multiplicity were introduced for model-based recursive partitioning by Zeileis and coworkers [9, 7]. Alternatively, permutation tests can be applied in situations where asymptotic normality of the partial score is not guaranteed [10] or in cases with small numbers of observations [11–13], which are common in medicine. Also in this case procedures for multiple testing are used to cope with a possibly large number of partitioning variables J .

If we can reject at least one of the $2 \times J$ null hypotheses for the global model $\mathcal{M}((Y, \mathbf{X}), \hat{\vartheta})$ at a pre-specified nominal level, model-based recursive partitioning selects the partitioning variable Z_{j^*} associated with the highest correlation to any of the partial score functions. This is usually done by means of the smallest p -value. The dependency structure between the partitioning variable Z_{j^*} and either one of the partial score functions is described by a simple cut-point model. Once we find an optimal cut-point $Z_{j^*} < \mu$ using a suitable criterion [13, 7], we split the patients into two subgroups according to $Z_{j^*} < \mu$. For both subgroups, we estimate two separate models with parameters $\hat{\vartheta}(1)$ and $\hat{\vartheta}(2)$, respectively, obtain the corresponding partial score functions, and test the independence hypotheses. If we find deviations from independence, we in turn estimate a cut-point in the most highly associated partitioning variable, and split again. The procedure of testing independence of partial score functions and partitioning variables is repeated recursively until deviations from independence can no longer be detected.

Since model-based recursive partitioning is a tree method, in the following we use topic-specific vocabulary, such as nodes. The root node contains all patients and is the basis for the initial model, inner nodes represent splits and leaf nodes contain the patients of the different subgroups and specify the partition-specific models. The paths from root to leaf nodes define the subgroups.

2.2 Content interpretation

A clearer picture of the interpretation of subgroup-dependent model parameters and distribution of the partial scores under unstable parameters is best given by means of a partitioned linear model discussed in the following.

Here x_A is a contrast that indicates whether a subject was treated with treatment A (active) but not C (control) in a two-armed trial and x_{stratum} is a stratum with $\mathbf{x} = (x_A, x_{\text{stratum}})$. The conditional distribution of the primary endpoints Y given treatment and stratum is normal

$$Y|\mathbf{X}=\mathbf{x} \sim \mathcal{N}(\alpha + \beta x_A + \gamma x_{\text{stratum}}, \sigma^2). \quad (4)$$

The segmented model we want to fit using model-based recursive partitioning reads

$$Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z} \sim \mathcal{N}(\alpha(\mathbf{z}) + \beta(\mathbf{z})x_A + \gamma x_{\text{stratum}}, \sigma^2), \quad (5)$$

where γ is the effect of the stratum and the variance σ^2 is a nuisance parameter. The objective function for a patient with observations $(y, \mathbf{x})$ is the negative log-likelihood, when maximum likelihood estimation is used, or the error sum of squares, when ordinary least squares is used. Yet, both methods lead to the same scores

$$\psi((y, \mathbf{x}), \hat{\theta}) = \left(\frac{\partial \Psi((y, \mathbf{x}), \theta)}{\partial \alpha} \Big|_{\theta=\hat{\theta}}, \frac{\partial \Psi((y, \mathbf{x}), \theta)}{\partial \beta} \Big|_{\theta=\hat{\theta}} \right)^T = \frac{1}{\sigma^2} \begin{pmatrix} y - (\hat{\alpha} + \hat{\beta}x_A + \hat{\gamma}x_{\text{stratum}}) \\ (y - (\hat{\alpha} + \hat{\beta}x_A + \hat{\gamma}x_{\text{stratum}})) \cdot x_A \end{pmatrix} \quad (6)$$

and thus to the same solution. Note that the partial score function with respect to the intercept is proportional to the least-square residuals and all further scores are proportional to the product of the residuals and the respective variable.

A partitioning variable can be predictive, prognostic, or both, and we have to consider the parameters in the model to understand the nature of a partitioning variable. Figure 1 shows examples for mean primary endpoints and the corresponding intercept α and treatment effect β . If $\alpha(\mathbf{z})$ varies over $\mathbf{z}$, but $\beta(\mathbf{z})$ is constant, then the components of $\mathbf{z}$ are prognostic because the mean primary endpoint varies but not the treatment effect (see first column of Figure 1). If $\beta(\mathbf{z})$ varies over $\mathbf{z}$ and $\alpha(\mathbf{z})$ is constant, then the variables in $\mathbf{z}$ are predictive since it means that the mean primary endpoint in one treatment arm stays the same but the treatment effect changes over $\mathbf{z}$ (second column). If both parameters vary, then $\mathbf{z}$ is predictive (third column) or predictive and prognostic at the same time (last column). In the latter situation, the mean primary endpoint of the second subgroup changes over $\mathbf{z}$ and the intercept also changes.

Figure 1: Possible mean primary endpoint within subgroups resulting from a predictive, prognostic, or predictive and prognostic variable.

It is also interesting to take a closer look at the partial scores. Figure 2a shows the partial scores with respect to intercept and treatment parameter that result from a linear model $Y|\mathbf{X}=\mathbf{x} \sim \mathcal{N}(\alpha + \beta x_A, \sigma^2)$ plotted against a partitioning variable z_1 , which is predictive and prognostic. The data-generating process of this model was suggested by Loh et al. [14] and is defined as

Figure 2: Partial scores of different kinds of variables. The symbols represent the treatment arms C and A as indicated. (a) Partial scores of a predictive and prognostic variable (eq. (7)). (b) Partial scores of a predictive and prognostic variable (eq. (8)). (c) Partial scores of a prognostic variable (eq. (9)).

$$Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z} \sim \mathcal{N}(1.9 + 0.2 \cdot x_A + 1.8 \cdot \mathbb{1}(z_1 < 0) + 3.6 \cdot \mathbb{1}(z_1 > 0) \cdot x_A, 0.7), \quad (7)$$

with X_A from $\mathcal{B}(1, 0.5)$ and Z_1 from $\mathcal{N}(0, 1)$. For the example, we used this process to draw a sample of 200 observations.

The partial scores with respect to the intercept ψ_α fluctuate randomly around zero over the whole range of z_1 . The partial scores with respect to the treatment parameter ψ_β change. Hence, in this situation, model-based recursive partitioning would detect a deviation from independence between ψ_β and z_1 and implement a split at approximately $z_1 < 0$. There is no chance of finding this cut-point by looking at the least-square residuals only, since a deviation of independence between ψ_α and z_1 is hardly visible in the scatterplot in

the left panel of Figure 2a. Figure 2b shows the partial scores obtained with a slightly modified data-generating process, where instead of $\mathbb{1}(z_1 > 0) \cdot x_A$, one has $\mathbb{1}(z_1 < 0) \cdot x_A$:

$$Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z} \sim \mathcal{N}(1.9 + 0.2 \cdot x_A + 1.8 \cdot \mathbb{1}(z_1 < 0) + 3.6 \cdot \mathbb{1}(z_1 < 0) \cdot x_A, 0.7). \quad (8)$$

Here the procedure would split the partial score with respect to the intercept, although z_1 is still prognostic and predictive at the same time.

If we focus on the prognostic variable z_1 in the model

$$Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z} \sim \mathcal{N}(2 \cdot x_A + \mathbb{1}(z_1 > 0), 0.7), \quad (9)$$

we see non-random patterns in both scores (see Figure 2c). Since the partial scores with respect to the treatment parameter are set to zero for treatment arm A, we would split on basis of the scores with respect to the intercept, just as a consequence of a higher power.

These three examples show that splitting in the partial score with respect to the intercept does not give any information about whether the partitioning variable is predictive or prognostic. It also does not make sense to choose to split only in the score with respect to the treatment parameter because one might miss important cut-points. In order to be able to say whether a partitioning variable is predictive or prognostic, it is not enough to know which partial scores are responsible for the split. It is necessary to consider the model parameters in the segmented model. If the treatment parameter β varies in the subgroups, then the chosen partitioning variables are predictive or both predictive and prognostic. If β is constant, the variables are only prognostic.

2.3 Relation to established procedures

Traditional approaches for subgroup identification are also based on a model for the primary endpoint, but the segmentation is implemented by means of varying coefficients. More precisely, the model includes interactions between treatment and the patient characteristics $\mathbf{z}$ in addition to the main effects

$$\begin{aligned} \mathbb{E}(Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z}) &= \alpha + \beta x_A + \gamma_{\text{prognostic}}^T \mathbf{z} + \gamma_{\text{predictive}}^T \mathbf{z} x_A \\ &= (\alpha_0 + \gamma_{0,z}^T \mathbf{z})(1 - x_A) + (\alpha_1 + \gamma_{1,z}^T \mathbf{z}) x_A, \end{aligned} \quad (10)$$

with $\alpha = \alpha_0$, $\beta = \alpha_1 - \alpha_0$, $\gamma_{\text{prognostic}}^T = \gamma_{0,z}^T$ and $\gamma_{\text{predictive}}^T = \gamma_{1,z}^T - \gamma_{0,z}^T$. The model is known as the “classical approach” for subgroup analyses [15, 16]. Significant interaction terms $\gamma_{\text{predictive}}$ are in this case subject to the choice of relevant partitioning variables. However, patient subgroups can only be identified directly in this model for categorical variables z_j since the model has no notion of optimal cut-off points. As the number of potential partitioning variables J might be large, the simultaneous estimation of all parameters in the model might be computationally burdensome and associated with a large variance. Regularisation procedures may be applied for selecting relevant interaction parameters that deviate considerably from zero.

RECPAM [17, 18] goes a step further and fits such models by trees. In every node, a likelihood-ratio test is computed that compares the segmented model

$$\mathbb{E}(Y|\mathbf{X}=\mathbf{x}, \mathbf{Z}=\mathbf{z}) = \alpha + \beta_1 x_A \mathbb{1}(z_j \in \mathcal{B}_k) + \beta_2 x_A [1 - \mathbb{1}(z_j \in \mathcal{B}_k)] \quad (11)$$

to the constant model

$$(\mathbb{E}Y|\mathbf{X}=\mathbf{x}) = \alpha + \beta x_A \quad (12)$$

for every possible segment $\mathcal{B}_k$ ($k=1, \dots, K$) induced by all possible cut-off points in z_j , i.e. an exhaustive search is performed. The procedure is applied to all partitioning variables z_j ($j=1, \dots, J$). The algorithm then chooses the variable and segmentation that comes along with the highest test statistic. The method is so far limited to linear models and Cox proportional hazards models, and parameter instabilities can only be detected in β but not in α .

A method that is similar in spirit to model-based recursive partitioning, but is limited to normal linear models, is the Gs method [14] based on the GUIDE algorithm [19, 20]. Instead of using partial scores with respect to intercept and treatment effect, Gs uses only the least-square residuals (that is, only the partial score with respect to the intercept). In contrast to model-based recursive partitioning, Gs looks at the dichotomised (at zero) residuals separately in the two treatment arms. The independency between positive/negative residual signs and each partitioning variable is tested using a chi-squared test separately for each treatment. If the partitioning variable is at least ordinal, it is dichotomised by splitting at the mean. The optimal split variable chosen is the one that induces the highest sum of chi-squared statistics. Looking at the left panels of Figures 2a and 2b, one can imagine that in these situations the procedure may successfully find the subgroups. However, in a less clear situation and where the optimal cut-point is not near the mean of z_1 , the method will have lower power or will not be able to find a split at all.

Another recently proposed tree algorithm is qualitative interaction trees (QUINT [21]). QUINT searches for instabilities in the treatment parameter β only, but the resulting partitions have to have different signs in the parameter. In other words, QUINT aims at finding subgroups in which the treatment effect is the reverse of that of the other subgroups. The current implementation of QUINT [22] is limited to continuous primary endpoints. It would be possible to enforce splits that are qualitatively different in model-based recursive partitioning. This could be achieved by incorporating a criterion that implements a split only if the treatment effects in the two new subgroups have different signs.

SIDES (subgroup identification based on differential effect search) [23] and SIDEScreen [24] aim at identifying subgroups of patients with high benefit from a novum compared to the standard treatment. Although the subgroups are linked to hypercubes in the sample space of Z , they are overlapping and can therefore not be represented as a tree structure. The methods are based on a cross-validated implementation of subgroups that were obtained on independent learning samples.

More general approaches blending recursive partitioning with traditional models (known as hybrid, model, or functional trees in machine learning Gama [25]), include M5 [26], GUIDE [19], CRUISE [27], LOTUS [28] and maximum likelihood trees [29]. (Bayesian approaches can be found in Chipman, George, and McCulloch [30]) and (Bernardo, Bayarri, Berger, Dawid, Heckerman, Smith, and West [31]). Except GUIDE, none of these methods has been studied in the specific context of subgroup analyses so far.

3 Partitioning effects of Riluzole on ALS patients

ALS is a neurodegenerative disease that causes weakness, muscle waste and paralysis. Currently the only drug on the market for treating ALS is Riluzole (Rilutek). It slows down disease progression but only modestly prolongs life expectancy by about two months [32]. A more thorough investigation of the treatment effect of Riluzole in ALS patients is of great importance since a cure is not yet available and patients usually die within 1.5–4 years after disease onset [33]. We use model-based recursive partitioning to address the question whether Riluzole has an especially low or high treatment effect on both functional and survival endpoints of any subgroups of patients.

Our analysis is based on patient information obtained from the PRO-ACT (Pooled Resource Open-Access ALS Clinical Trials) database [34], which contains data of ALS patients that were involved in one of several publicly- and privately-conducted clinical trials. The database provides information on patient survival, functional endpoint (the ALS functional rating scale), Riluzole use, demographics, family history, patient history, forced and slow vital capacity, laboratory data and vital signs. The data were fully de-identified and therefore the centres of data ascertainment are not given in the data set. The participants gave their informed consent, and study protocols were approved in the respective medical centres. The database was initiated by the non-profit organisation Prize4Life that aims at accelerating cure and drug development for ALS, for example through the DREAM-Phil Bowen ALS Prediction Prize4Life challenge [35].

The ALS Functional Rating Scale (ALSFRS [36]), is a widely used instrument for evaluating the functional status of patients with ALS even though the uni-dimensionality of the score seems questionable [37]. It is a sum-score of the following ten items: speech, salivation, swallowing, handwriting, cutting food and handling utensils, dressing and hygiene, turning in bed and adjusting bed clothes, walking, climbing stairs, and breathing. Each of these items can have values from zero to four, where four is normal and zero indicates the inability of performing the respective action. Hence, if the ALSFRS has a value 40, the patient has normal abilities for all items. The lower the score, the worse is the patient's status. The items were measured at several time points during the study period. We focused on the ALSFRS reading approximately six months after treatment start as the functional endpoint. Approximately means that we used the measurement closest to six months after treatment start, with a maximal absolute deviation of 20 days. In addition, we also decomposed the score and modelled the items defining the score separately.

The survival time of patients was measured in days starting with the patient's enrolment in one of the trials. For patients without survival information, we used the latest follow-up time given for the patient in the data as censoring time.

Model-based recursive partitioning was applied to models for the functional and survival endpoints. We allowed parameter instabilities in both the intercept and the Riluzole treatment effect. Bonferroni-adjusted permutation tests using test statistics of a quadratic form [13] were applied for assessing independence of the partial score functions and the partitioning variables and also for cut-point selection. The use of permutation tests for cut-point selection improves speed compared to the original suggestion of fitting and comparing models for all reasonable partitions [7]. We restricted the depth of the trees to two levels. Parameter estimates including confidence intervals are given for the final subgroups. Note that we are computing the confidence intervals after applying model selection through splitting into subgroups and thus the intervals should be interpreted with caution. For both endpoints, we used partitioning variables available at patient enrolment from the following groups of variables: demographics, family history, patient history, forced and slow vital capacity, laboratory data, and vital signs. We excluded patient records with missing values at the endpoints; the sample size was $N = 2534$ for the functional endpoint and $N = 3306$ with 916 events for the survival endpoint.

3.1 ALSFRS

The ALSFRS six months after treatment start (ALSFRS_6) defined the functional endpoint. The sum-score is positive, and the model needs to adjust for the baseline ALSFRS obtained at treatment start (ALSFRS_0). We used a normal GLM with log-link and offset $\log(\text{ALSFRS}_0)$ such that the model

$$\mathbb{E}\left(\frac{\text{ALSFRS}_6}{\text{ALSFRS}_0} \mid X = x\right) = \frac{\mathbb{E}(\text{ALSFRS}_6 \mid X = x)}{\text{ALSFRS}_0} = \exp\{\alpha + \beta x_R\} \quad (13)$$

describes the expected relative change in the ALSFRS over the first six months under treatment. The treatment (Riluzole/no Riluzole) is indicated by x_R . The model was fitted by maximum likelihood.

The time between disease onset and start of treatment, the forced vital capacity (FVC), and the phosphorus balance are the three partitioning variables selected for the tree given in Figure 3. The FVC value gives the volume of air in liters that can forcibly be blown out after full inspiration to the lung. A normal phosphorus balance is between 1 and 1.5 mmol/L. The tree indicates a negative treatment effect of Riluzole for patients with fewer days between disease onset and start of treatment that have a higher FVC value (node 4). Therefore, the FVC value is predictive in the group of patients with less than 468 days between disease onset and treatment start. Patients with more days between disease onset and treatment start do not seem to have a treatment effect. The fact that the time since onset plays an important role is not surprising since it is a surrogate for the speed of disease progression [38]. Patients with a slow progression were seldom included early in one of the studies. Hence a long time between onset and start of treatment usually stands for a slow progression.

Figure 3: Results of application of model-based recursive partitioning with a Gaussian GLM with log link and offset on the data from the PRO-ACT database with the ALSFRS score as primary endpoint variable. Inner nodes give the split variable selected and the associated permutation test based p -value for the split. Terminal nodes give the model coefficients including standard confidence intervals.

3.2 ALSFRS items

The model for the ALSFRS sum-score assumes that the effect of Riluzole is the same for the ten items that define the score. In a more fine-grained analysis, we decomposed the score into its ten items (each ranging between zero and four) and modelled each item by means of a proportional odds model. For one of the ten items assessed at six months, e.g. Y_6 , the model reads

$$\mathbb{P}(Y_6 \leq r | \mathbf{X} = \mathbf{x}) = \frac{1}{1 + \exp(-\alpha_r + \beta x_R)}, \quad (14)$$

where $r = 0, \dots, 4$ is one of the five possible values of Y . The intercept parameters are now $\alpha = (\alpha_0, \dots, \alpha_3)$ and the partial score function ψ_α is now four dimensional.

As in the previous example, we needed to adjust for the baseline value Y_0 , i.e. the value of the ALSFRS item read at the beginning of treatment. This adjustment was implemented by computing separate models; one each for the observations with a start value k , which allows a baseline-specific intercept and treatment effect :

$$\mathbb{P}(Y_6 \leq r | Y_0 = k, X = x) = \frac{1}{1 + \exp(-\alpha_{rk} + \beta_k x_R)} \quad \text{for } k = 0, \dots, 4. \quad (15)$$

Therefore, we had a total of five different treatment parameters and 20 different intercepts for each of the ten different items. Model-based recursive partitioning was used to assess the parameter instability of all 250 parameters simultaneously. Note that some of these parameters could not be estimated owing to too small of sample sizes; these were simply discarded.

The implementation of the non-standard model in the theoretical and computational framework of model-based recursive partitioning was straightforward. For every node, we computed the five separate models for the respective baseline values for each of the ten items and extracted the partial scores. A stratified permutation test using the baseline values as independent blocks was used to assess parameter instability. The same procedure was applied for cut-off selection.

The resulting tree (on top of Table 1) contains splits in time between disease onset and treatment start and in the FVC value. The tree is in good agreement with the tree based on the ALSFRS (Figure 3). The third split variable is the lymphocyte percentage. Normal lymphocyte concentrations range from 16 to 33%. Table 1 shows the coefficient values of the models in the terminal nodes for every item and every starting value of the given item. Empty fields indicate that it was not possible to compute the model. Obviously, there were not enough observations in models with zero as starting value for any items in any nodes. The colours in the table indicate whether the effect of Riluzole was positive (blue), negative (pink) or zero (grey). The colours were assigned on the basis of confidence intervals of the coefficient in the given model. Riluzole had a positive effect on patients in the partition of terminal node 3 who had a starting value of 4 in item 1 (speech), 3 (swallowing) or 9 (climbing stairs) and on patients in the partition of terminal node 7 that had a starting value of 3 in item 5 (cutting food and handling utensils). Patients in node 4 who had a starting value of 3 in item 6 (dressing and hygiene) had a negative effect of Riluzole. Riluzole had no effect on patients in the partition of node 6 which are the patients with more than 584 days between disease onset and treatment start who have a lymphocyte concentration under 21.5%.

3.3 Survival time

We used both a Weibull model and a Cox model to identify subgroups with differing effects of Riluzole on the survival endpoint. The application of the model-based recursive partitioning framework in the Weibull model is straightforward and was introduced by Zeileis et al. [7]. Since the Cox model is a semiparametric model, where the intercept is a function of time, treated as a nuisance parameter omitted in the partial likelihood, there is no direct way of obtaining ψ_α . Because, conceptually, deviance residuals are always

Table 1: Coefficient and confidence interval of Riluzole use in the terminal nodes for every item and every starting value in the model-based recursive partitioning with a proportional odds model (ALSFRS items as outcome). Blue indicates a positive effect of Riluzole, pink a negative effect and grey no effect.

Item	No.	Start	Node 3	Node 4	Node 6	Node 7
Speech	1	0				
		1				
		2		−0.27 (−1.15, 0.60)		
		3	0.33 (−0.33, 1.00)	0.04 (−0.40, 0.47)	−0.27 (−1.12, 0.56)	−0.06 (−0.60, 0.48)
		4	0.84 (0.08, 1.59)	0.11 (−0.28, 0.48)		
Salivation	2	0				
		1				
		2	0.22 (−0.94, 1.40)	0.43 (−0.48, 1.36)	1.36 (−0.31, 3.11)	−0.20 (−1.28, 0.88)
		3	0.15 (−0.57, 0.87)	−0.26 (−0.76, 0.23)	−0.24 (−0.96, 0.47)	−0.05 (−0.58, 0.48)
		4	0.49 (−0.11, 1.07)	−0.03 (−0.39, 0.32)		
Swallowing	3	0				
		1				
		2	0.35 (−0.62, 1.32)	−0.89 (−2.06, 0.21)	1.51 (−0.33, 3.51)	−0.75 (−1.96, 0.43)
		3	0.57 (−0.06, 1.21)	−0.36 (−0.85, 0.12)	−0.40 (−1.17, 0.36)	0.35 (−0.22, 0.93)
		4	0.62 (0.01, 1.23)		0.15 (−0.49, 0.75)	0.28 (−0.22, 0.75)
Handwriting	4	0				−1.45 (−3.21, 0.37)
		1		−1.15 (−2.54, 0.20)	−0.36 (−1.93, 1.25)	−0.08 (−1.26, 1.12)
		2		−0.54 (−1.33, 0.25)		0.04 (−0.71, 0.79)
		3	−0.13 (−0.78, 0.51)	0.14 (−0.22, 0.49)	−0.08 (−0.70, 0.52)	−0.28 (−0.74, 0.17)
		4	−0.10 (−0.71, 0.49)	0.04 (−0.34, 0.42)		−0.14 (−0.65, 0.36)
Cutting	5	0				
		1			−0.01 (−1.19, 1.15)	−0.79 (−1.68, 0.07)
		2		0.15 (−0.54, 0.85)		0.48 (−0.14, 1.12)
		3	−0.03 (−0.76, 0.70)	−0.07 (−0.44, 0.30)	0.10 (−0.60, 0.79)	0.52 (0.03, 1.02)
		4	0.13 (−0.45, 0.72)	−0.09 (−0.49, 0.31)	−0.14 (−0.82, 0.53)	−0.21 (−0.74, 0.30)
Hygiene	6	0				
		1				
		2		−0.11 (−0.65, 0.43)		−0.37 (−0.89, 0.15)
		3	−0.22 (−0.88, 0.44)	−0.37 (−0.72, −0.03)	0.14 (−0.50, 0.78)	0.27 (−0.17, 0.71)
		4	0.26 (−0.40, 0.92)	0.01 (−0.42, 0.44)	0.14 (−0.71, 0.98)	0.30 (−0.31, 0.90)
Bed	7	0				
		1				
		2			−0.03 (−0.96, 0.89)	0.29 (−0.47, 1.04)
		3	0.15 (−0.57, 0.87)	−0.32 (−0.71, 0.08)	−0.12 (−0.74, 0.49)	−0.05 (−0.45, 0.35)
		4	−0.21 (−0.80, 0.36)	−0.10 (−0.45, 0.24)	−0.11 (−0.81, 0.57)	−0.35 (−0.90, 0.18)

(continued)

Table 1: (continued)

Item	No.	Start	Node 3	Node 4	Node 6	Node 7
Walking	8	0				
		1				
		2			0.48 (−0.22, 1.16)	−0.04 (−0.56, 0.46)
		3		0.11 (−0.33, 0.55)		0.46 (−0.12, 1.04)
		4	0.51 (−0.16, 1.18)	0.13 (−0.27, 0.52)		
Stairs	9	0				
		1		−0.02 (−0.49, 0.45)	−0.01 (−0.72, 0.68)	−0.39 (−0.89, 0.10)
		2			−0.80 (−2.10, 0.48)	0.07 (−0.97, 1.12)
		3		0.26 (−0.19, 0.72)	−0.65 (−1.46, 0.15)	−0.16 (−0.79, 0.46)
		4	1.01 (0.27, 1.77)	0.06 (−0.35, 0.48)	0.72 (−0.11, 1.55)	0.29 (−0.32, 0.89)
Respiratory	10	0				
		1				
		2				
		3				
		4	0.58 (−0.06, 1.24)	−0.08 (−0.44, 0.28)		

defined as the derivative of the log-likelihood with respect to the intercept, we applied martingale residuals as ψ_α . Also worth noting is that both models assume proportional hazards. For the segmented model, proportional hazards are only assumed within each partition. This has to be kept in mind when interpreting the treatment effect in different nodes: Parameters with different signs are clearly linked to opposing treatment effects, but when the parameters only differ in size, it is hard to say whether it is because the groups differ in treatment effect or because they differ in the hazard function.

3.3.1 Weibull model

The Weibull model is a transformation model of the form

$$\mathbb{P}(Y \leq y | X = x) = F\left(\frac{\log(y) - \alpha_1 - \beta x_R}{\alpha_2}\right), \quad (16)$$

where F is the cumulative distribution function of the Gompertz distribution. Weibull models are fitted via maximum-likelihood estimation, and therefore the objective function in this case is the negative log-likelihood and the score function has one column per parameter, i.e. intercept α_1 , slope parameter β and scale parameter α_2 . In the Weibull model, we take the usual intercept as well as the scale parameter as “intercept”-parameter $\alpha = (\alpha_1, \alpha_2)^T$ because they define the shape of the baseline hazard and hence in some respect take the role of an intercept. Splitting in the intercept or scale parameter score suggests non-proportional hazards.

Figure 4 shows that the patient’s age and again the time between onset and treatment start play a role in the partitioning. Older patients (> 55.7 years) for whom the time between onset and treatment was longer than 757 days and very young patients did not seem to benefit at all from the treatment. In the remaining two groups, life expectancy seemed to be prolonged for patients treated with Riluzole.

3.3.2 Cox model

The use of the Cox model in model-based recursive partitioning is a rather special case, since the baseline hazard in the Cox model is treated as an infinite-dimensional nuisance parameter and estimation is performed by minimisation of the negative partial log-likelihood. The Cox proportional hazards model is given by

Figure 4: Results of application of model-based recursive partitioning with a Weibull model and data from the PRO-ACT database with survival time as primary endpoint variable. Inner nodes give the split variable selected and the associated permutation test based p -value for the split. Terminal nodes give the model coefficients, including standard confidence intervals and the survival curves in the two groups of treatment. Rugs indicate event times.

$$\lambda(y|\mathbf{x}) = \lambda_0(y) \exp(\beta x_R), \quad (17)$$

where λ is the hazard function and λ_0 the baseline hazard function. The partial score function ψ_α (or better, ψ_{λ_0}) cannot be easily derived. As surrogate score function, we propose using the martingale residuals as a score for the baseline hazard, which takes the role of an intercept in the Cox model, and the score residuals for the treatment parameters β . The score residuals are an intuitive choice because they are the first derivative of the partial log-likelihood with respect to the parameters. We used martingale residuals to check whether there is a general difference in the endpoint for different patients, which in parametric models is usually shown by the score with respect to the intercept. Instability in the martingale residuals indicates a violation of the proportional hazards assumption. Since the martingale residuals are not normally distributed, the application of permutation tests is more appropriate than the use of M-fluctuation tests.

Age and the time between disease onset and start of Riluzole treatment form the segments in this example. The tree in this example has almost the same splits as the tree in the previous example. Also estimates support the results of the Weibull example. Again, we did not see much difference between treated and untreated very young patients. For all other groups, Riluzole treatment led to a slight tendency for a lower risk of death (Figure 5).

4 Discussion

Model-based recursive partitioning allows the direct segmentation of the model describing the overall treatment effect as specified in the study protocol. This is the most important benefit of embedding subgroup analysis into this framework because it would be hard to explain why the overall treatment effect and the partitioned treatment effect have to be estimated by two different procedures. This renders the application of suboptimal models unnecessary, such as when a change score is analysed using linear models [21].

Although we are conceptually only interested in finding predictive factors, we think it is necessary to allow splits in the partial scores with respect to both intercept and treatment parameter. This procedure will also detect prognostic factors, but there is a higher chance of including all relevant predictive factors since one might miss prognostic factors when only the treatment scores are split. In our analysis, we decided on the nature of the partitioning variables (prognostic or predictive) only when we interpreted the results of the analysis.

In a model with more covariates than the treatment (e.g. strata), we would still split the partial scores with respect to intercept and treatment parameter for subgroup analyses. A theoretical assumption is then that the parameters that are not split stay constant. In practice, this assumption usually does not hold. It is generally also possible to split more than just the scores with respect to intercept and treatment parameter. Then the split variables are not restricted to being predictive or prognostic but may have an association with the effect of the other covariates.

In model-based recursive partitioning, the variable selection in each node is error controlled, i.e. the probability of selecting a partitioning variable for splitting, when actually all partitioning variables are independent of the scores, is at most as large as the nominal level. The only drawback of using multiple testing procedures is in cases where there are many possible partitioning variables that do not contain information, because with increasing number of noise variables the chance of detecting an actually existing subgroup goes down. The application of permutation tests has the advantage of taking the correlation structure among the partitioning variables into account. Furthermore, for small studies or small subgroups, the exact conditional p -value can be easily approximated up to any desired accuracy; therefore, the method does not rely on asymptotic arguments. The trees obtained by model-based recursive partitioning allow straightforward visualisation, potentially enriched with plots illustrating the distribution of the endpoints

Figure 5: Results of application of model-based recursive partitioning with a Cox model and data from the PRO-ACT database with the survival time as primary endpoint variable. Inner nodes give the split variable selected and the associated permutation test based p -value for the split. Terminal nodes give the model coefficients including confidence intervals and the survival curves in the two groups of treatment. Rugs indicate event times.

for the different treatment groups in each subgroup. Therefore, the results of such a subgroup analysis are easily communicated to physicians. Looking at a tree is much easier than trying to understand the meaning of higher-order interactions in a linear predictor. A general drawback of tree methods is the instability of the tree structure with respect to small perturbations in the data, whereas the resulting partitions we are primarily interested in are often relatively stable [13]. Instability in the tree structure can be assessed by means of the variable selection and split statistics, where it is easy to identify all equally likely splits. Bootstrap aggregation and forest procedures are well-known for their ability to stabilise single trees [39] at the cost of interpretability and point into a promising future research direction also for model-based recursive partitioning.

The statistical properties of the confidence intervals derived from the segmented model await further attention. Leeb and Pötscher ([40]) discuss the validity of inference after variable selection and claim that it is difficult if at all possible. Bai and Perron [41], who discuss the construction of confidence intervals after splitting up the data based on a break point in a single partitioning variable, argue that it is possible. In our approach we first search for the most appropriate partitioning variable (variable selection) and then search for the optimal split point (break point selection). To our knowledge there is no literature on inference after variable and break point selection and thus it is unclear if or how valid confidence intervals can be computed. In any case the results of such a subgroup analysis have to be confirmed in follow-up trials, which lowers the necessity of confidence intervals. To be conservative one can see the confidence intervals for parameters in the subgroup-specific models as shown in our examples as a range of possible values and hence as a measure of variability rather than significance [42].

It would be interesting to extend the framework of the PRO-ACT database of ALS studies to models for non-independent data, such as mixed models for longitudinal observations. This would allow ALS disease progression to be modelled over time, and also a potentially time-varying treatment effect to be assessed. In our way of modelling the functional endpoint, we include no information about patients that died within the first six months after treatment start. Joint modelling of the longitudinal functional endpoint and the survival endpoint is a means of combining all possible information [43].

Despite the deficits of model-based recursive partitioning for subgroup analysis discussed in this section, we think that the procedure as introduced and illustrated in this paper rather closely resembles the requirements for statistical procedures in this field as outlined in the EMA guideline [1]. In particular, it is the most generally applicable procedure with statistical error control and unbiased variable selection [13, 7]. With the available open-source implementation (see following section for details), the method can be applied straightforwardly elsewhere.

Computational details

An open-source implementation of all methods discussed in this paper and beyond is available in the **partykit** package `hothorn_partykit_2014`. PRO-ACT data are available at <https://nctu.partners.org/ProACT/> [44]. The source code for reading and cleaning the database is provided in the **TH.data** package [45]. The source code for the analyses is provided in the supplementary material. All computations were conducted using **partykit** (version 0.8-2) in the R system for statistical computing [46], version 3.1.2).

Listing 1: Code snippet for Weibull model in model-based recursive partitioning using the function `ctree()` from the **partykit** package.

```
## Function to compute Weibull model and return score matrix
mywb <- function(data, weights, parm) {
  mod <- survreg(Surv(survival.time, cens) Riluzole,
    data = data, subset = weights > 0,
    dist = "weibull")
}
```

```

ef <- as.matrix(estfun(mod) [,parm])
ret <- matrix(0, nrow = nrow(data), ncol = ncol(ef))
ret[weights > 0,] <- ef
ret
}
## Compute tree
tree <- ctree(fm, data = data, ytrafo = my.wb,
              control = ctree_control(maxdepth = 2,
                                     testtype = "Bonferroni"))

```

Acknowledgements: We are thankful to the organisers and participants of the “Workshop on Classification and Regression Trees” (March 2014), sponsored by the Institute for Mathematical Sciences of the National University of Singapore, for helpful feedback and stimulating discussions and to Karen A. Brune for improving the language. Financial support by the Swiss National Science Foundation (grant 205321_163456) is gratefully acknowledged.

References

1. European Medicines Agency. EMA guideline on the investigation of subgroups in confirmatory clinical trials (draft). Available at: http://www.ema.europa.eu/docs/en_GB/document_library/Scientific_guideline/2014/02/WC500160523.pdf, 2014.
2. Gadbury GL, Iyer HK. UnitTreatment Interaction and Its Practical Consequences. *Biometrics* 2000;56:882–5.
3. Holland PW. Statistics and causal inference. *J Am Stat Assoc* 1986;81:945–60.
4. Morgan JN, Sonquist JA. Problems in the analysis of survey data, and a proposal. *J Am Stat Assoc* 1963;58:415–34.
5. Doove LL, Dusseldorp E, Van Deun K, Van Mechelen I. A comparison of five recursive partitioning methods to find person subgroups involved in meaningful treatment-subgroup interactions. *Adv Data Anal Class* 2014;8:403–25.
6. Italiano A. Prognostic or predictive? It's time to get back to definitions!. *J Clin Oncol* 2011;29:4718–4718.
7. Zeileis A, Hothorn T, Hornik K. Model-based recursive partitioning. *J Comput Graph Stat* 2008;17:492–514.
8. Hastie T, Tibshirani R. Varying-Coefficient Models. *J R Stat Soc Series B (Methodological)* 1993;55:757–96. Available at: <http://www.jstor.org/stable/2345993>.
9. Zeileis A, Hornik K. Generalized M-fluctuation tests for parameter instability. *Stat Neerl* 2007;61:488–508.
10. Zeileis A, Hothorn T. A toolbox of permutation tests for structural change. *Stat Papers* 2013;54:931–54.
11. Hothorn T, Hornik K, Van de Wiel MA, Zeileis A. A Lego system for conditional inference. *Am Stat* 2006a;60:257–63.
12. Hothorn T, Hornik K, van de Wiel MA, Zeileis A. Implementing a class of permutation tests: The coin package. *J Stat Software* 2008;28:1–23.
13. Hothorn T, Hornik K, Zeileis A. Unbiased recursive partitioning: A conditional inference framework. *J Comput Graph Stat* 2006b;15:651–74.
14. Loh W-Y, He X, Man M. A regression tree approach to identifying subgroups with differential treatment effects. *Stat Med* 2015;34:1818–33.
15. Foster JC, Taylor JM, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Stat Med* 2011;30:2867–80.
16. Kehl V, Ulm K. Responder identification in clinical trials with censored data. *Comput Stat Data Anal* 2006;50:1338–55.
17. Ciampi A, Negassa A, Lou Z. Tree-structured prediction for censored survival-data and the Cox model. *J Clin Epidemiol* 1995;48:675–89.
18. Negassa A, Ciampi A, Abrahamowicz M, Shapiro S, Boivin JF. Tree-structured subgroup analysis for censored survival data: Validation of computationally inexpensive model selection criteria. *Stat Comput* 2005;15:231–9.
19. Loh W-Y. Regression trees with unbiased variable selection and interaction detection. *Stat Sinica* 2002;12:361–86.
20. Loh W-Y. Improving the precision of classification trees. *Ann Appl Stat* 2009;3:1710–37.
21. Dusseldorp E, Van Mechelen I. Qualitative interaction trees: A tool to identify qualitative treatment-subgroup interactions. *Stat Med* 2013;33:219–37.
22. Dusseldorp E, Doove L, Van Mechelen I. quint: Qualitative Interaction Trees. Available at: <http://CRAN.R-project.org/package=quint>, R package version 1.0., 2013
23. Lipkovich I, Dmitrienko A, Denne J, Enas G. Subgroup identification based on differential effect search – A recursive partitioning method for establishing response to treatment in patient subpopulations. *Stat Med* 2011;30:2601–21.

24. Lipkovich I, Itri AD. Strategies for identifying predictive biomarkers and subgroups with enhanced treatment effect in clinical trials using SIDES. *J Biopharm Stat* 2014;24:130–53.
25. Gama J. Functional trees. *Mach Learn* 2004;55:219–50.
26. Quinlan JR. C4.5: Programs for Machine Learning. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc, 1993.
27. Kim H, Loh W-Y. Classification trees with unbiased multiway splits. *J Am Stat Assoc* 2001;96:589–604.
28. Chan K-Y, Loh W-Y. Lotus: An algorithm for building accurate and comprehensible logistic regression trees. *J Comput Graph Stat* 2004;13:826–52.
29. Su X, Wang M, Fan J. Maximum likelihood regression trees. *J Comput Graph Stat* 2004;13:586–98.
30. Chipman H, George E, McCulloch R. Bayesian treed models. *Mach Learn* 2002;48:299–320.
31. Bernardo JM, Bayarri MJ, Berger JO, Dawid AP, Heckerman D, Smith AFM, et al. Bayesian Treed Generalized Linear Models. *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, 85, 2003.
32. European Medicines Agency. Riluzole Zentiva: EPAR summary for the public. Available at: http://www.ema.europa.eu/docs/en_GB/document_library/EPAR_-_Summary_for_the_public/human/002622/WC500127609.pdf, 2012.
33. Chiò A, Logroscino G, Hardiman O, Swinger R, Mitchell D, Beghi E, et al., and On Behalf of the Eurals Consortium. Prognostic factors in ALS: A critical review. *Amyotroph Lateral Scler* 2009;10:310–23.
34. Atassi N, Berry J, Shui A, Zach N, Sherman A, Sinani E, et al. The PROACT database: Design, initial analyses, and predictive features. *Neurology* 2014;83:1719–25.
35. Küffner R, Zach N, Norel R, Hawe J, Schoenfeld D, Wang L, et al. Crowd sourced analysis of clinical trial data to predict amyotrophic lateral sclerosis progression. *Nat Biotechnol* 2014;33:51–7.
36. Brooks BR, Sanjak M, Ringel S, England J, Brinkmann J, Pestronk A, et al. The amyotrophic lateral sclerosis functional rating scale – assessment of activities of daily living in patients with amyotrophic lateral sclerosis. *Arch Neurol* 1996;53:141–7.
37. Franchignoni F, Mora G, Giordano A, Volanti P, Chio A. Evidence of multidimensionality in the ALSFRS-R scale: A critical appraisal on its measurement properties using Rasch analysis. *J Neurol Neurosurg Psychiatry* 2013;84:1340–5.
38. Hothorn T, Jung HH. RandomForest4life: A Random Forest for predicting ALS disease progression. *Amyotroph Lateral Scler Frontotemporal Degener* 2014;15:444–52.
39. Strobl C, Malley J, Tutz G. An Introduction to Recursive Partitioning: Rationale, Application and Characteristics of Classification and Regression Trees, Bagging and Random Forests. *Psychol Methods* 2009;14:323–48.
40. Leeb H, Pötscher BM. Model selection and inference: Facts and fiction. *Economet Theory* 2005;21:21–59.
41. Bai J, Perron P. Computation and analysis of multiple structural change models. *J Appl Economet* 2003;18:1–22.
42. Lagakos SW. The challenge of subgroup analyses—reporting without distorting. *N Eng J Med* 2006;354:1667–9.
43. Henderson R, Diggle P, Dobson A. Joint modelling of longitudinal measurements and event time data. *Biostatistics* 2000;1:465–80.
44. Massachusetts General Hospital, N. C. R. I. Pooled resource open-access ALS clinical trials database. Available at: <https://nctu.partners.org/ProACT/>, 2013.
45. Hothorn T. TH.data: TH's Data Archive. Available at: <http://CRAN.R-project.org/package=TH.data>, R package version 1.0-4, 2014.
46. R Core Team. R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria. Available at: <http://www.R-project.org/> 2014.
47. Hothorn T, Zeileis A. partykit: A modular toolkit for recursive partitioning in R, Technical report. Available at: <http://eecon.uibk.ac.at/wopec2/repec/inn/wpaper/2014-10.pdf>, 2014.